Highlights

CogniTrade: A Multi-Agent Framework for Data-Driven Financial Decision Making Using Fuzzy Inference and Genetic Programming

Srikari Rallabandi,Manleen Kaur,Pranav Misra

- Innovative Paradigm for Financial Forecasting:

  This study introduces a groundbreaking decision support system that revolutionizes stock selection through a novel combination of multi-agent modeling, fuzzy inference systems, and genetic programming

  The framework harnesses the power of collaborative agent networks and fuzzy logic systems to decipher optimal trading directions and strengths, providing deep insights into the complex behavior of financial markets.

- Multi-Agent System for Enhanced Decision Making:

  A zero-sum multi-agent system based on fuzzy inference systems and genetic programming is designed to predict directional movements and magnitudes of stock prices.

  Agents within the community collaborate to generate trading signals by aggregating technical indicators and producing directional strength values for future market behavior.

- Dynamic Adaptation for Improved Performance:

  The study introduces a dynamic adaptation strategy for the replacement chance parameter, ensuring that the genetic programming process avoids stagnation and continues to evolve more effective solutions.

  A fuzzy inference system assesses the stagnation index and adjusts the replace chance parameter accordingly, leading to better convergence and enhanced forecasting accuracy.

- Integration of Technical Indicators and Fuzzy Logic:

  The proposed method seamlessly integrates technical indicators (ADX, SO, RSI) and fuzzy logic to create agent-based models that predict directional strengths of stock movements.

  Fuzzy inference systems utilize evolved membership functions to interpret and aggregate input signals, producing actionable predictions for traders.

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

CogniTrade: A Multi-Agent Framework for Data-Driven Financial Decision Making Using Fuzzy Inference and Genetic Programming

Srikari Rallabandi[a], Manleen Kaur[b] and Pranav Misra[c]

[a] *Vidya Jyothi Institute of Technology, Hyderabad, India*

[b] *Kalinga University, Raipur, India*

[c] *Dr. A.P.J Abdul Kalam Technical University, Prayagraj, India*

ARTICLE INFO

Keywords:

Artificial intelligence

Financial forecasting

Fuzzy Systems

Genetic Programming

Multi-Agent Systems

Abstract

This paper introduces a pioneering decision support system that revolutionizes stock selection through an innovative blend of a multi-agent paradigm, fuzzy inference systems, and genetic programming. In this framework, a collaborative network of agents harnesses fuzzy logic systems to decipher optimal trading directions and strengths, shedding light on the intricate nuances of financial market behavior. To arrive at a decision, an architecture of a zero-sum multi-agent system based on fuzzy inference systems and genetic programming was designed. The system adeptly predicts directional movements and magnitudes by conducting regression tasks on stock time series data. The multi-agent system performs a regression task on the time series data of a stock to determine the direction and strength of the next movement in the time series; By adopting this framework, individuals can examine its affiliation functions to comprehend financial market behavior better. Rigorous testing against a Dow Jones Index benchmark underscores the method's exceptional profitability, outperforming established techniques. With its potential to amplify investment decisions through insightful data analysis, this system emerges as a transformative force in financial decision-making.

1. Introduction

Many decades ago, financial traders implemented statistical models to forecast time series, creating a field called technical analysis. This type of market analysis is characterized by using past financial data to describe certain aspects of a market, such as its volatility, trend, and momentum. These tools receive the name technical indicators (TI). Some examples of TI are those created by Welles Wilder, such as the relative strength index and average directional index Wilder (1978). An example of an oscillator, a type of TI that usually tells if a market is overbought or oversold (the market has been up-trending or down-trending, respectively, for a relatively long time) is the stochastic oscillator Schirding (1984), created by George Lane. Unlike technical analysis, which primarily depends on price movements in a time series, fundamental analysis comprehensively evaluates the intrinsic factors influencing a specific financial market or multiple markets. This approach focuses on understanding the deeper forces that shape the market's dynamics rather than merely relying on historical price data. A technique that can be used in fundamental analysis is to select a group of companies and analyze their economic behaviors to establish a general forecast for each. After performing this analysis, one can narrow the group of companies and decide on a proper trading strategy to follow according to the arrived conclusions. To analyze a company, one can examine its business plan and management.

A more robust approach to performing these analyses is machine learning techniques. With machine learning, one can use inputs provided by technical or fundamental indicators and create regression models. Examples of regression techniques are autoregression Burg (1968), symbolic regression Billard and Diday (2002), and linear regression Kutner, Nachtsheim and Neter (2004). Several advanced methodologies are available in machine learning to address regression

sreeku.ralla@gmail.com (S. Rallabandi); manlinsarna74@gmail.com (M. Kaur); pranavmisra7@gmail.com (P. Misra)
https://www.linkedin.com/in/srikari-r-ab2829173/ (S. Rallabandi)
ORCID(s):

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

or curve-fitting assignments, including utilizing hybrid artificial neural networks Melin, Mancilla, Lopez and Mendoza (2007) and the application of support vector regression Basak, Pal and Patranabis (2007). But a problem that often arises with these models is that they can't account for emergent phenomena: they learn from past behaviors and can't predict what they haven't "seen." An approach that has shown excellent results in simulating these behaviors is agent-based modeling (ABM). Citing the work by Bonabeau Bonabeau (2002), ABM is "by its very nature, the canonical approach to modeling emergent phenomena," and for this reason, it is a suitable tool to model complex systems Jennings (2001), such as financial markets. According to Gilbert's work Gilbert (2008), ABM can be succinctly defined as a computational approach within the realm of social science. It involves devising an abstraction of "social reality," aiming to portray one's perception of how reality functions vividly. In particular, ABM has been proven to be useful in analyzing price stability in financial markets Pellizzari and Forno (2007), scientometrics and domain visualization Niazi and Hussain (2011), and in social sciences in general Gilbert (2008), among others.

This work proposes a method based on ABM to construct a decision support system (DSS) that aids traders on the process of trading financial markets. DSS is not a well-defined term, as for some researchers is just an interactive system for managers. For others, the focus is on understanding and improving the decision process Keen (1980). Since its inception, DSS has been created to aid managers in business-related areas Sprague (1980) Little (1979), and one can see specialized cases where DSS are used to support the decision process when forecasting financial markets, such as in the work by Tsang, Yung, and Li Tsang, Yung and Li (2004). The proposed method aims to improve the decision process of trading financial markets: the trader receives a recommendation of what financial market to trade, as in the work by Brown, Pelosi, and Dirska Brown, Pelosi and Dirska (2013). But the created algorithm goes deeper than just giving recommendations; in the end, a novel method for performing regression was created, and the authors believe that it can be extended to perform classification tasks. The method uses a combination of fuzzy inference systems (FIS), genetic programming (GP) Poli, Langdon, McPhee and Koza (2008) Koza (1992), and multi-agent systems (MAS) Shoham and Leyton-brown (2009) to produce models that can explain the behavior of complex systems such as that of a financial market by analyzing the evolved membership functions (MF) of each agent.

The differences between ABM and MAS are subtle but must be noted before proceeding. According to Niazi and Hussain Niazi and Hussain (2011), MAS is a sub-domain of ABM; the principal objective of ABM is to explain a phenomenon through the interaction of agents, while MAS provides a specific application of ABM to solve a practical problem. The MAS developed in this work belongs to a class of coalitional games called constant-sum game Shoham and Leyton-brown (2009). Here, agents are selfish and don't intervene with others, but a higher-order process dictates how they must organize. In this case, the collective work of every agent must give a "zero-sum": the sum of their forces must be equal to the observed price, or in other words, they must adjust to the real prices. To achieve this, each agent uses a FIS to decide how much trading force it will provide and in what direction, and its MF is generated using GP. The combination of using FIS and GP allows us to create conclusions about the behavior of the market by interpreting the evolved MF of each agent in the system (for example, one can infer that whenever certain TI is low and another TI is high, the market will stagnate). Specifically, the predicates and consequents of the FIS are generated using the sum of sines found using a modified algorithm based on GP. Sum of sines was used because, in theory, they can fit any curve. Finally, the resulting outcome of the devised algorithm is the agent community that achieves the minimum mean-squared error (MSE) in absolute terms, combined with the summation of their individual scores (elaborated in Section 3).

The organization of this investigation is delineated in the following manner: Section 2 is comprehensive in examining relevant prior research. This section was given special consideration to establish the essential groundwork for the proposed approach, which incorporates diverse technologies and methodologies. Section 3 delves into the theoretical underpinnings of the suggested methodology, thoroughly explaining its specifications. Section 4 presents a series of extensive experiments and their outcomes. The experiments reflect thoughtful design considerations when furnishing substantial evidence to support the effectiveness of the devised technique. Lastly, Section 5 offers concluding remarks about the presented work and briefly outlines future research endeavors.

2. Related Work

In Subsection 2.1, we begin by discussing various studies related to market prediction and the diverse methods employed to analyze and anticipate financial markets proficiently. Several and different approaches of the design of MAS are presented in a series of works in Subsection 2.2. Methods related to the proposed method that used evolutionary

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

algorithms to optimize or find better architectures in the development of forecasting models are presented in Subsections 2.3 and 2.4. Considering the studies highlighted in Subsection 2.4 is crucial since the devised approach specifically incorporates GP as a vital component of its algorithm. The next Subsection 2.5 presents works where Fuzzy Logic is used to forecast financial markets. Finally, some works about DSS are briefly mentioned in Subsection 2.6.

2.1. Financial Forecasting

There is a plethora of techniques used for regression and the creation of DSS, and it would be inconvenient to list many of them, but some notable examples are presented.

Brown, Pelosi, and Dirska Brown et al. (2013) apply a specialized genetic dynamic-radius species-conserving algorithm to identify Dow Jones Index stocks for investment. It is important to mention this work because, in the end, the DSS presented in the proposed method makes the same recommendation as in their work. Additionally, Brown, et al., uploaded the dataset that the authors of the present work used to perform the different experiments. In Section 4, a comparison to their work is provided.

Lu, Lee, and Chiu Lu, Lee and Chiu (2009) underscore the intricacy of financial time series. They emphasize its turbulent nature and propose a methodology to address this turbulence, employing a two-stage modeling process that integrates dimensionality reduction and regression analysis. Their technique initially applies dimensionality reduction to generate uncorrelated components, identifying and filtering out those affected by disturbances. Subsequently, the remaining components are harnessed to reassemble the variables of interest; they encompass diminished disturbances, acting like input for the regression analysis model. Their work significantly influenced the development of the suggested approach, as we posit that the ABM methodology can be utilized to mitigate market disturbances, leveraging a distinct group of agents dedicated to its modeling.

Finally, it is crucial to acknowledge the application of neural architecture in regression tasks, as it has been demonstrated as an exceptionally efficient technique for solving such problems. O'Connor and Madden Connor and Madden (2005) obtained some remarkable results, obtaining an annual 23.5% investment rate (ROI) on Dow Jones data for training and testing. Another example is given by Castillo and Melin Castillo and Melin (2001), where they assess diverse hybrid structures that integrate deep learning and rule-based systems for forecasting stock market trends.

2.2. Multi Agent Systems

A multi-agent system is the fundamental algorithm at its highest level for the proposed method. It is, therefore, paramount to mention some related works that use MAS to forecast or understand financial markets.

The research by Klingert and Meyer Klingert and Meyer (2012) employs a MAS to investigate the impact of continuous trading auction and logarithmic price scoring mechanisms. Utilizing an agent-based simulation model, they aim to evaluate how these mechanisms influence trading activity, prediction market precision, and price variability. The researchers validate three proposed hypotheses, and ultimately, their findings suggest that the logarithmic price scoring mechanism demonstrates a competitive advantage over the alternative mechanism, driven by increased trades and enhanced price stability.

Sherstov and Stone Sherstov and Stone (2005) present three automated stock-trading agents who follow different strategies to predict financial markets and are compared. The first agent uses reinforcement learning, the second a trend-following strategy, and the last one is market-making. These agents are part of a MAS where the better-performing agent is chosen for the testing phase.

In their investigation, Kendall and Su Kendall and Su (2003) employed a MAS to construct a simulation of financial markets. The simulation encompassed representing stock traders as diverse, adaptable synthetic actors. Interestingly, approximately 80% of these fabricated equity traders proficiently executed trading tactics, leading to exceptional gains for investors when juxtaposed with an uncomplicated buy-and-hold approach.

The authors of this work gained useful knowledge about MAS from two theses. The first one is the work from Grothmann Grothmann (2002). This work served as inspiration for the architecture of the proposed method. The second thesis is by Boer-Sorbán Boer-Sorban (2008), which gave an overview of approaches to describe and understand financial market's dynamics, and motivated the authors of this work to use the approach of agent-based computation to perform financial forecast.

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

As a final mention, Samanidou et al. Samanidou, Zschischang, Stauffer and Lux (2007), provide the reader a comprehensive overview of ABM, where different techniques to perform this kind of model are discussed.

2.3. Genetic Algorithms

In the proposal, GP-utilization generates an MF of FIS that act as the agents' functions. What follows is the mention of two works that use genetic algorithms to perform such a task, and in Subsection 2.4, one can find more specialized works where GP is used.

Thrift Thrift (1991) explores a nowadays widely used genetic algorithm-based technique to discover MF parameters in a FIS to obtain better performance. Homaifar and McCormick Homaifar and McCormick (1995) explore a more comprehensive genetic algorithmic technique to optimize the MF alongside the rules for fuzzy logic controllers.

2.4. Genetic Programming

GP is commonly used to create regression models or find a better architecture in trading strategies. In the context of this research, an unconventional technique is adopted in which GP is employed to generate the MF of the FIS, serving as the agents' functions within the MAS. As far as the authors know, this approach is not commonly utilized. What follows is a set of works on some common uses of GP in financial applications.

Li and Tsang Li and Tsang (1999) developed a system that generates decision trees of TI using GP. Initial findings demonstrate its superior performance compared to widely employed generic trading strategies concerning forecast precision and average yearly compounded profits during two separate out-of-sample test durations, each lasting three and a half years.

Garcia-Almanza and Tsang Garcia-Almanza and Tsang (2006) used GP as a regression tool and a repository method technique to model "rare instances" or emergent phenomena. The repository technique examines decision trees derived through GP to identify categorization guidelines for infrequent cases. This method allows for diverse modeling of exceptional instances, thereby increasing the likelihood of recognizing analogous scenarios in future time series data. The insights from Garcia-Almanza's research were beneficial for the authors of this work, as it emphasized the significance of identifying these emerging phenomena in financial markets and highlighted the utility of ABM in capturing and modeling such occurrences.

An artificial market that models technical, fundamental, and noise traders was developed by Martinez-Jaramillo and Tsang Martinez-Jaramillo and Tsang (2009). In this work, GP is used to generate the agent functions Russell, Norvig, Canny, Malik and Edwards (2003) of the technical traders, and TI was used as the operator for the GP algorithm.

Chen and Yeh Chen and Yeh (2001) introduce a novel approach that integrates GP and ABM, designed to incorporate trader's (agents) exploration patterns. Utilizing simulated annealing, this framework establishes correlations between these patterns and social influences, such as group dynamics or macroeconomic conditions like the standard of living.

Utilizing GP, Bastian Bastian (2000) deduces the input parameters, establishes the rule framework, and models the membership characteristics of fuzzy systems.

2.5. Fuzzy Prediction

As has been previously mentioned, in this work, each agent uses a FIS as its agent's function. The agent function aims to determine how much "trading force" an agent has to provide, so the sum of all the agent forces gives a zero-sum system. FIS has successfully been used to predict financial markets in the past.

Ijegwa et al. Ijegwa, Rebecca, Olusegun and Isaac (2014) developed a FIS that utilizes four TI as inputs to generate signals for action, indicating whether to "go long," "go short," or "stay neutral" in the financial market. This system is a valuable decision support tool, assisting traders in their analysis and strategy formulation. The selected technical indicators encompass the exponential moving average convergence/divergence, relative strength index, stochastic oscillator, and accumulation/distribution volume.

Huang, Pasquier, and Quek Huang, Pasquier and Quek (2009) delve into utilizing an advanced coevolutionary fuzzy system named HiCEFS in financial time series forecasting. Their methodology revolves around the TI known as the relative strength index (RSI), a widely accepted tool for analysis. An intriguing aspect of their investigation lies in

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

incorporating irregular shaped membership functions (ISMF). By integrating ISMF, their system constructs a more precise predictive tool, which exhibits superior performance compared to a simplistic passive investment approach.

Gamil, et al. Gamil, El-fouly and Darwish (2007) devised a resilient framework that synergizes MAS, fuzzy logic, genetic algorithms, and moving averages as vital components of the FIS. Integrating these elements into the system allows it to interpret and utilize stock information from various sources effectively. This collective approach streamlines the prediction process for investment decisions, as the computational demands necessitate enhanced processing capabilities. The system optimizes decision-making and augments overall efficiency by employing diverse agents for intelligent guidance.

2.6. Decision Support Systems

This final subsection presents two works related to DSS and financial forecast, focusing on DSS development and deployment.

First up is the work by Keen Keen (1980), which helped the authors of the present work to understand DSS. His work presents valuable definitions related to DSS and design guidelines. Numerous empirical investigations illustrate crafting methodologies pertaining to DSS.

Secondly, Tsang, et al. Tsang, Li and Butler (1998) Tsang, Li, Markose, Er, Salhi and Iori (2000) Tsang et al. (2004) Tsang (2008), present a DSS, over a series of works, for financial forecasting called EDDIE. The system aims to enhance a trader's likelihood of executing successful trades, acting as a dynamic decision support mechanism. Its performance depends on the accuracy and relevance of user data, as well as the effectiveness of its GP exploration algorithm.

2.7. Latest Trends

- Hybrid Models: Many recent works have focused on combining different machine learning techniques to create hybrid models for financial forecasting.Nguyen, Nguyen and Nguyen (2019) These models often blend the strengths of various algorithmsHussain, Muhammad and Khan (2019) to improve accuracy and robustness.
- Deep Learning: Deep learning techniques, especially neural networks, have gained popularity in financial forecasting due to their ability to capture complex relationships in data.Alam, Lee and Egiazarian (2018) L. Fang, et al. have explored various architectures, such as LSTM and CNN, to improve prediction accuracy.Fang, Wang, Guo, Zhang and Wei (2021)
- Ensemble Methods: Ensemble methods, which combine predictions from multiple models, have been widely used to enhance forecasting performance.Ferreira, da Costa, Rodrigues and Papa (2019) These methods aim to reduce the impact of individual model biases and errors.
- Explainability: With the increasing adoption of machine learning in finance, there's a growing emphasis on interpretability and explainability.Dash and Behera (2020) Researchers are developing techniques to make complex models more transparent and understandable.

3. Proposed Method

3.1. Dataset

The dataset used in this work was first used by Brown, Pelosi, and Dirska Brown et al. (2013), and can be found at the UCI Machine Learning Repository under the name of "Dow Jones Index Data Set." This research exclusively employed the market's opening, peak, trough, and closing values as our primary dataset. To obtain additional indicators, namely the average directional index (ADX), stochastic oscillator (SO), and relative strength index (RSI), we employed an online tool called FreeStockCharts, which is accessible at http://www.freestockcharts.com/. We added TI to all stocks provided in the dataset by Brown et al. using FreeStockCharts and downloaded the full historical data from the platform. Filtering was performed where only the relevant records (the same ones in the benchmarking dataset) were extracted.

As a final step, another attribute was added to the dataset, which the authors of this work call "directional strength." This attribute tells a movement's strength at a particular time in a time series and its direction (downtrend or uptrend). An absolute value of 100 of the directional strength means that it was a very strong movement, where all the traders in that financial market agree to buy or sell, and a value of 0 means strong indecision. The directional strength indicator could

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

have already been mentioned in the literature, but the authors of this work couldn't find a reference to this particular formula. The formula to calculate the directional strength is:

$$DS(t) = 100\frac{C(t) - O(t)}{H(t) - L(t)} \tag{1}$$

t represents a particular point of time in a time series, O represents the open price, H the high, L the low, and C the close price.

3.2. Decision Support System

The final outcome of the proposed method is a recommendation to go long, go short, or hold in financial markets. Users are told which direction the market will take (downtrend or uptrend) and its strength through the use of the directional strength indicator (see Subsection 3.1).

Nevertheless, the underlying method can be used to perform regression tasks, and it could be extended to perform classification tasks in future work. What follows is the explanation of this underlying method.

3.3. Communities of Agents

As has been noted before, the proposed method is based on ABM. In ABM implementation, we developed a conceptual community of agents. A community is a collection of agents who are the final actuators in the environment (the simulated financial market). The action of each of the agents in the community involves providing a certain trading force, which is a value between -100 and 100. The sum of all of the forces of the agents in a community must be equal to the observed or real directional strength indicator for a particular record in the dataset. A community receives its inputs, which are the TI ADX(t), SO(t), and RSI(t), and gives as output the DS(t+1) (the directional strength of a future record, i.e. one data point in the future).

The method creates several communities, which exchange the genetic material of its agents among them, in order to produce better-performing communities, and their fitness is determined using a sum of MSE and the absolute value of the sum of the forces of its agents. This summation was determined to be a good fitness indicator by trial and error, but other fitness functions can be used. A flowchart depicting the algorithm that creates a population of these communities is presented in Figure 1.

3.4. Fuzzy Inference Systems

Each community agent has a FIS that acts as its agent function. The inputs of this FIS, as has been mentioned before, are the ADX, SO, and RSI. The output ranges from -100 to 100, capturing directional strength. To arrive to this output number, one can choose a number of MF that act as inputs and a number of MF that act as outputs. By trial and error, the authors found that a number of 5 outputs gave good results in most cases. The FIS fuzzified the aggregation of the MF at their activation level determined by its inputs, representing a directional strength for that particular agent.

The rule set is depicted below and is comprised of 15 rules: 5 outputs for each of the 3 TI. GPpred$N$ represents a predicate MF created by the GP algorithm and, similarly, GPconsN represents a consequent.

1. if ADX is GPpred1 then DS is GPcons1
2. if ADX is GPpred2 then DS is GPcons2
3. if ADX is GPpred3 then DS is GPcons3
4. if ADX is GPpred4 then DS is GPcons4
5. if ADX is GPpred5 then DS is GPcons5
6. if RSI is GPpred6 then DS is GPcons6
7. if RSI is GPpred7 then DS is GPcons7
8. if RSI is GPpred8 then DS is GPcons8
9. if RSI is GPpred9 then DS is GPcons9
10. if RSI is GPpred10 then DS is GPcons10
11. if SO is GPpred11 then DS is GPcons11

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

12. if SO is GPpred12 then DS is GPcons12
13. if SO is GPpred13 then DS is GPcons13
14. if SO is GPpred14 then DS is GPcons14
15. if SO is GPpred15 then DS is GPcons15

3.4.1. Genetic Programming

To create the MF, a GP algorithm performs a number of operations among the initially generated communities. These operations are the classical crossover and mutation operations and two new operations: migration and replacement. The crossover operation chooses a subtree of the GP function of the *i*th agent of a community A and passes it to the GP function of the *i*th agent of a Community B. A subtree from the community B is passed to the GP function of the *i*th agent. As a result, two new agents are generated and replace the old agents. To determine what communities are to be crossed over, a tournament is held where two random communities are chosen from the population, and the one with the lowest MSE plus the absolute value of the sum of forces (as explained in Subsection 3.3 is chosen; this will obtain the first community. This process is repeated to obtain a second community. As for the mutation operation, a subtree from community A is selected and randomly changed to produce a new GP function. The community to be mutated is randomly selected from the population.

The migration operation is a rather simple one. According to a migration chance, an agent from community A will get interchanged with another agent from community B. To determine if this interchange is to be performed, the algorithm checks if the sum of the forces of community A is of a different sign than the sum of forces of community B. This ensures that the migration will produce a community closer to a zero-sum, i.e. the community will be closer to the observed or real prices. As for the replacement operation, a whole new community is generated and replaces the community with the highest absolute value of the MSE plus the sum of forces. All four operations have a pre-established chance of occurring (except for the replace chance, which is currently being dynamically adapted by using a FIS. See Subsection 3.5). One can see a flowchart of the general process in Figure 1.

In addition to this set of operations, the GP algorithm uses the parameters of population size, community-wise tally of agents, total generation count, the number of inputs processed by agents (in this case, the TI), and how many outputs (which, in the end, as has been noted previously, get defuzzified into a directional strength value).

How the proposed method generates the MF is a novel or, at least uncommon method. The GP algorithm uses a sum of sines to produce functions converted to the MF of the FIS of the agents. A graphical representation of an input MF and an output MF can be seen in Figures 3 4, respectively. The GP algorithm uses only the sum and the sine operators. The sum of sines is represented by Formula 2.

$$y = \sum_{i=1}^{n} a_i sin(b_i x + c_i) \tag{2}$$

Taking into consideration Formula 2, the GP algorithm uses the following literals set: a, b, and c are random real numbers, where a ranges from 0.0 to 0.3, b ranges from 0.0 to 10.0, and c ranges from 0.0 to 1.0. These ranges were obtained by trial and error, and other ranges could perform better. An additional real number literal is added to the set, which ranges from -1 to 1, along with the variable x, which takes the value of the current input (ADX, RSI or SO).

3.5. Dynamic Adaptation of the Replace Chance

The replace chance parameter, instead of being constant, is dynamically adapted using a FIS while the GP process is running. This idea was borrowed from the series of works by Castillo, et al. An example of this dynamic adaptation of parameters can be found in Castillo, Neyoy, Soria, Melin and Valdez (2015), where the authors apply a FIS that controls the value of certain parameters in an ant colony optimization algorithm (ACO), proposing different architectures. Then, these architectures are compared in performance. Ultimately, the authors demonstrate that employing FIS for dynamic parameter adaptation can enhance the effectiveness of ACO. Additionally, one can find a survey of the application of this type of technique in the work by Valdez, Melin and Castillo Valdez, Melin and Castillo (2014).

In the proposed method, a similar approach is taken, where dynamic adaptation of parameters through FIS is used to avoid the stagnation of the GP algorithm. In this case, a stagnated GP algorithm would mean that it has been producing

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

the same fitness over a number of generations, i.e., the absolute value of MSE plus the sum of forces is not being lowered in the last N generations. To calculate this stagnation level, the authors of this work propose the stagnation index, which is described by Formula 3, where S(P) means the stagnation index of P, where P is the set of fitnesses over a certain number of generations, and C(P) is the count of repeated fitnesses over the last N generations. For example, given the set P = 5, 5, 5, 3, 2, 1, C(P) = 3.

$$S(P) = 1 - \sqrt{\frac{1}{C(P)}} \tag{3}$$

The stagnation index is given as input to a FIS, which gives as output a new replace chance for the replace operation. This dynamic adaptation of parameters aims to obtain lower errors in a smaller number of generations. This is proven in Section 4. Also, it is worth mentioning that this technique could be applied to the crossover, migration, and mutation chances. Still, due to a time limitation, the authors of this work could only apply the technique to one parameter.

Figure 2 displays a flowchart of the GP technique devised.

3.6. How does our work build on recent developments

Our work aligns with the latest trends mentioned in the section 2.7 and contributes to advancing the state of the art:

- Hybrid Approach: Our proposed method leverages a hybrid approach by combining fuzzy inference systems, genetic programming, and multi-agent systems. This approach allows your model to capture both the interpretability of fuzzy logic and the complexity of genetic programming, enhancing the accuracy of financial forecasting.
- Novelty in Agent-Based Modeling: While many recent works have used machine learning techniques, your approach introduces agent-based modeling (ABM) to financial forecasting. This novel integration of ABM adds a layer of understanding to market dynamics that traditional models might miss.
- Emergence and Self-Organization: Our work capitalizes on the emergent behaviors and self-organization inherent in multi-agent systems. By allowing agents to interact and adapt to market conditions, your method captures the dynamic nature of financial markets, aligning with recent trends in capturing emergent phenomena.
- Interpretable Results: Unlike some complex machine learning models, your method produces interpretable results through evolved membership functions. This contributes to the trend of explainability in financial modeling, enabling traders and analysts to understand the rationale behind predictions.

4. Experimentation

The section commences by contrasting the novel approach with the method introduced by Brown, Pelosi, and Dirska Brown et al. (2013), employing their specified dataset to facilitate a thorough evaluation. To augment the statistical validity, the original experiment devised by Brown et al. was expanded, comprising 30 testing iterations instead of the initial 8. Additionally, an extensive array of experiments was conducted, incorporating a progressive increase in generational iterations, with each experiment comprising 30 trials, closely mirroring Brown et al.'s methodology in selecting a winning stock. Subsequently, the section provides a comprehensive exposition of the experiment and its outcomes achieved by implementing the FIS-guided adaptive replacement strategy. Furthermore, an additional set of experiments was conducted involving various configurations of the proposed method, and each configuration was executed over 30 trials to enable the collection and in-depth analysis of the resultant average return on investment (ROI).

Following the initial set of trials' conclusion, the researchers deliberately chose to embark on a distinct series of experiments, centering on the individual forecasting of stocks instead of the weekly stock selection approach. This novel set of experiments adopts a regression-oriented methodology rather than a recommendation-based one. Initially, a careful selection of five stocks was performed, applied considering various parameter setups. Subsequently, the performance evaluation of each stock in the dataset was executed using the proposed approach, applying a specific set of parameters that had demonstrated promising results after empirical testing. These experiments' effectiveness is assessed based on the collective average return on investment (ROI) achieved.

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

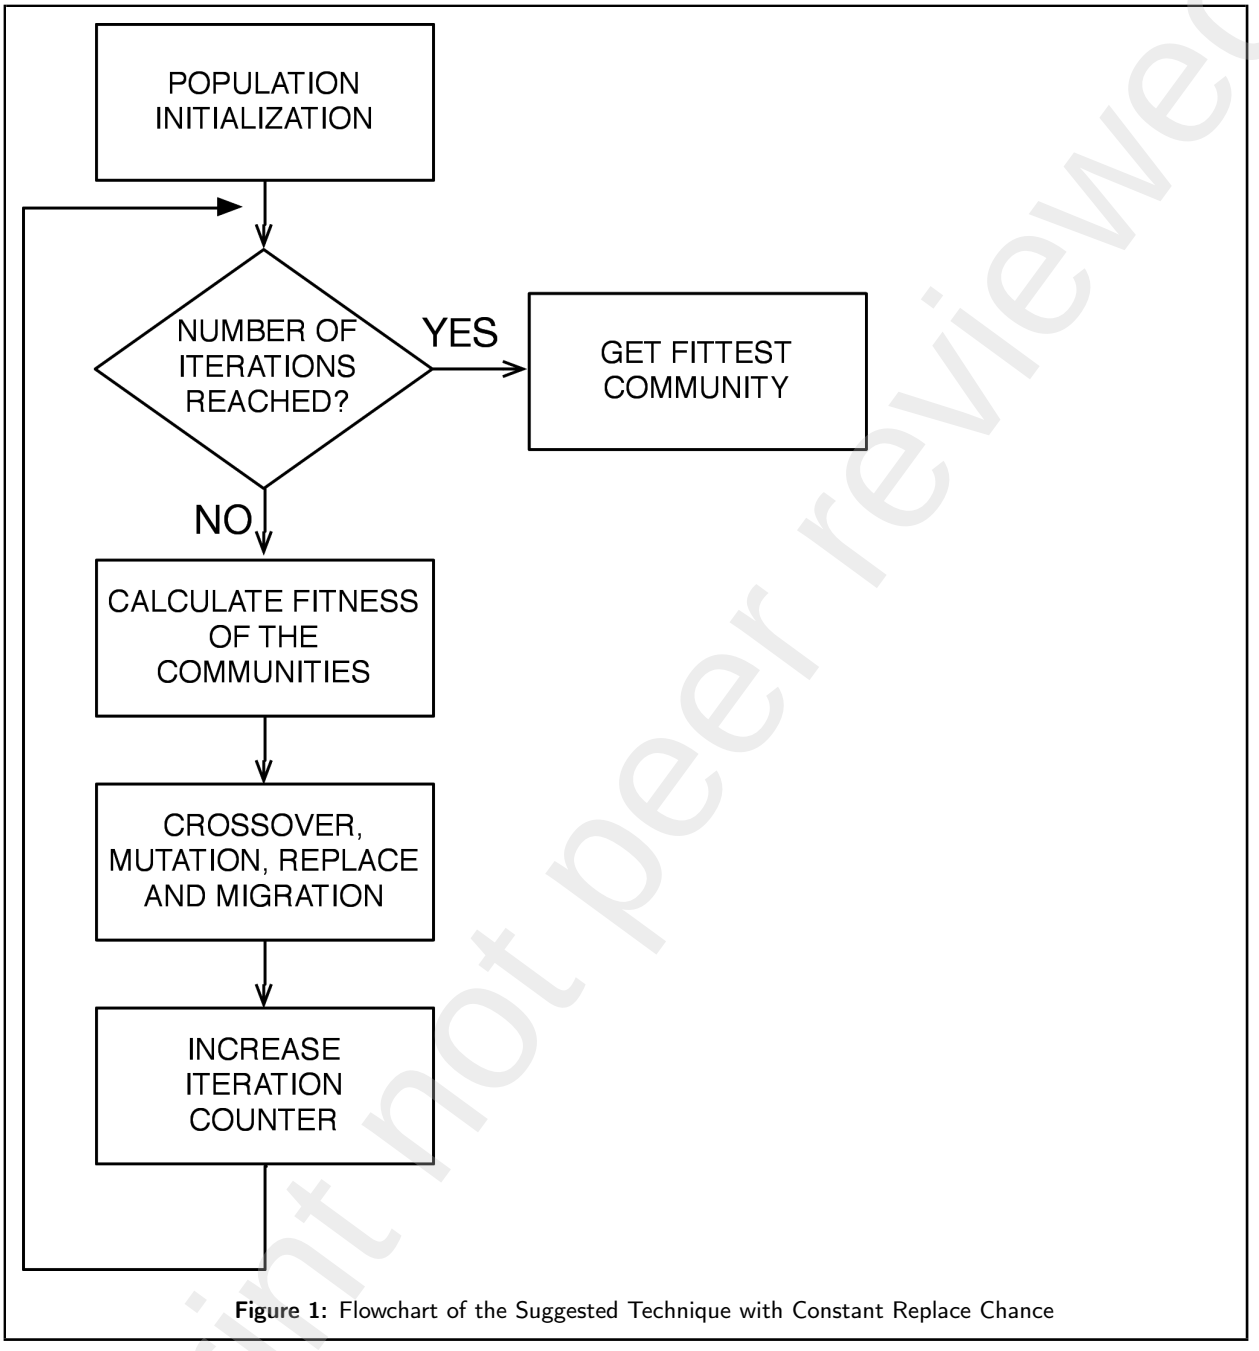

Figure 1: Flowchart of the Suggested Technique with Constant Replace Chance

As a concluding trial, the suggested approach undertakes a genuine regression endeavor to anticipate the directional potency of next week's entry in the collected data.

4.1. A Dynamic Approach for Evolving Replacement Probability

To evaluate the effectiveness of our suggested technique that included and excluded replacement chance's continuous adaptability, we conducted a couple of sets of iterations and then ran a hypothesis test for comparison. We conducted 30 experiments in each set, each running for 1000 generations with specific configurations. The first set employed a fixed replace chance, while the second set utilized a FIS to control the chance; refer to sub-section 3.5.

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Table 1
Average Errors - Dynamic vs. Fixed Mutation

Without Dynamic Adaptation		With Dynamic Adaptation	
1311.36858	1046.104578	1216.510814	1143.389484
1229.470821	1355.991986	1241.307708	1102.711613
1295.573386	1091.08208	1120.805705	1249.783318
1030.289883	1327.723752	1172.208267	1111.148651
1336.959564	1091.20479	972.4929395	1215.476516
1384.057403	1265.682395	1170.073256	1169.588612
1303.139391	1246.964085	1278.311636	1040.093053
1391.94702	1158.730397	1154.370559	762.2223289
1125.607147	1373.704464	1063.93828	1207.359942
1336.603628	1130.282865	1227.164144	1168.572853
1308.989252	1135.924961	1067.38689	1162.467748
1253.008383	1384.15463	1273.537452	1248.3565
1265.489682	1487.437018	960.6685944	1140.389075
1247.161813	1215.320663	1173.268224	1309.74542
1025.353754	1173.167377	1169.425716	978.8573602

Throughout each experiment, we recorded the error for every generation and subsequently calculated the average error across all generations upon completion of the experiment. This process was repeated for all 30 experiments in each set, resulting in 60 experiments. Table 1 displays the averaged error values obtained from these experiments.

The FIS design is straightforward, taking the stagnation index as input and giving as output a novel replacement chance in each run. Figure 14 illustrates the input membership functions (MFs), while Figure 15 displays the output MFs. The authors subjectively determined the design of these MFs through trial and error. Here is the set of rules:

1. should the SI be Low then the RC is also Low
2. should the SI be Medium then the RC is also Medium
3. should the SI be High then the RC is also High

Figure 16 demonstrates the effectiveness of dynamic adaptation of the replace chance variable in negating tendencies to stagnate. To rigorously test this hypothesis, we conducted a hypothesis test using Stata, and the results are displayed in Figure 17. The resulting t-score was -3.35, with a 95% confidence interval t-score of -2.0018. The null hypothesis is rejected, signifying a significant difference in the average error between experiments with dynamic adaptation and those without. The proposed method with dynamic adaptation outperforms the constant replace chance approach.

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Table 2
Stocks Selection for Trial 4

Week	Stock	Return		Week	Stock	Return
1	MMM	0.58265		8	MMM	1.27858
2	HD	1.92256		9	IBM	-2.01259
3	AA	3.72861		10	INTC	-1.92661
4	CSCO	3.48494		11	PFE	0.896414
5	CSCO	0.285551		12	AA	3.81731
6	HD	0.216626		13	PFE	3.23383
7	HD	0.981194		Total		16.489065

4.2. Selecting Stocks

Throughout the following experiments, we made use of the dataset supplied in Brown et al. (2013). The training phase involved the first quarter (12 records) of each stock, while the testing phase incorporated the second quarter (13 records).

4.2.1. How we fared vs. Brown et al.

In the research conducted by Brown and colleagues Brown et al. (2013), they employed their method to conduct a series of 8 trials with the objective of determining the best-performing stock to invest in for the upcoming week. Each trial involved assessing the rate of return achieved from purchasing all the stocks recommended by the algorithm. The authors presented a comprehensive table displaying the chosen stocks and the total rate of return obtained from their best trial, which happened to be the 4th trial in their experiments. To further illustrate the results, they used a chart resembling Figure 18, presenting the outcomes from all eight trials. Additionally, they provided an average of the results obtained across these multiple trials.

The authors of this research acknowledge that the current experiment alone does not provide enough evidence to fully demonstrate the effectiveness of their proposed method. Nevertheless, they offer an experiment akin to the one conducted by Brown et al. for the purpose of comparison. Table 2 displays the stocks chosen by their method along with their respective total rate of return. Notably, the results obtained by their approach outperformed those reported by Brown et al.

As previously indicated, Figure 18 presents a graphical depiction of the outcomes derived from the 8 experiments, portraying the quarterly gain and the mean weekly gain. In this instance, the mean rate of return achieved through these 8 experiments stood at 8.1140%, surpassing once more the mean rate reported by Brown et al., which amounted to 7.075%.

4.2.2. Expanded Trials - Unvarying Setup, Escalating Iterations

The extended experiments involved applying the proposed technique with a distinct setup comprising 4 clusters, 4 agents, 3 inputs, and 5 outputs. This setup was identified through iterative testing and demonstrated promising outcomes. We conducted 30 trials using this configuration, each with a varying number of generations. The outcomes were then averaged, and the results are presented in Figure 19.

The graph illustrates the method's behavior under different generations. When using a low number of generations, the approach tends to be undertrained, leading to suboptimal performance. As the number of generations increases to a moderate value, the method begins to yield improved results. However, as the number of generations further increases, the approach becomes overtrained, resulting in diminishing returns. With a very high number of generations, the method starts to exhibit chaotic behavior. Notably, throughout the experiments, the proposed method never produced a negative average rate of return.

Figure 5 displays a heatmap that allows the reader to understand the achieved results more deeply.

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Table 3
Choosing what Stock to Buy - Variable Configurations

	NUMBER OF AGENTS				
GEN	1	2	5	10	20
2	1.27%	2.46%	2.33%	1.68%	2.44%
5	3.22%	2.57%	0.97%	2.90%	3.70%
10	1.82%	1.75%	3.04%	3.19%	3.14%
20	2.32%	1.77%	2.28%	1.76%	4.82%
40	3.35%	4.07%	0.58%	4.13%	2.53%
80	1.48%	0.85%	3.34%	5.61%	2.81%

Table 4
American Airlines Rates of Return

	TRAINING SET SIZE			
GEN	2	4	6	8
5	-5.77% (6)	6.50% (8)	9.57% (9)	12.30% (8)
10	-5.77% (6)	6.50% (8)	7.68% (8)	12.30% (8)
15	-5.77% (6)	6.50% (8)	7.68% (8)	12.30% (8)
30	-5.77% (6)	14.95% (9)	7.68% (8)	9.13% (8)

4.2.3. Expanded Trials - Variable Settings

In this extended experiment, we follow a similar approach to that presented by Brown, et al., with some notable differences. Instead of conducting only 8 trials, we expanded the scope and ran a total of 30 trials. Furthermore, we explored various configurations of the proposed method. The results of these trials are summarized in Table 3.

4.2.4. Enhanced Trials - Diverse Configurations, Forecasting Individual Stocks

Tables 3, 5, 6, 7, and 8 present the results of the experiments, including the achieved rate of return and the number of correct predictions (shown in parentheses) for each configuration. Throughout this series of experiments, the main focus was on predicting the appropriate action (buy or sell) for a single selected stock, rather than stock selection. To complement the tabular data, visual representations of the most successful configurations are provided in Figures 13, 6, 7, 8, and 9.

4.2.5. Expanded Trials - Unvarying Setup, Forecasting for All Stocks

Table 9 presents the results of experiments conducted for each of the 30 stocks available in the dataset. For these trials, we used 12 weeks for training, followed by the next 13 weeks for testing. The objective of the proposed method was to predict whether the trader should buy or sell each stock. Depending on the accuracy of the predictions, we accumulated positive or negative rates of return. This process was repeated 30 times for each stock, and the outcomes are summarized in the table.

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Table 5
CISCO Rates of Return

	TRAINING SET SIZE			
GEN	2	4	6	8
5	15.38% (8)	26.60% (10)	14.42% (10)	12.62% (10)
10	15.38% (8)	26.60% (10)	14.42% (10)	19.62% (10)
15	15.38% (8)	26.60% (10)	13.46% (9)	15.15% (9)
30	13.02% (7)	26.60% (10)	13.46% (9)	19.62% (10)

Table 6
IBM Rates of Return

	TRAINING SET SIZE			
GEN	2	4	6	8
5	-10.43% (5)	-12.17% (5)	1.65% (9)	-4.16% (6)
10	-5.76% (6)	-4.18% (7)	-1.85% (8)	-2.29% (6)
15	-9.98% (6)	-7.21% (5)	-6.69% (7)	-2.65% (5)
30	-4.58% (6)	-9.61% (6)	2.99% (8)	-1.10% (7)

Table 7
The Coca-Cola Co. Rates of Return

	TRAINING SET SIZE			
GEN	2	4	6	8
5	-8.02% (5)	1.86% (6)	2.19% (7)	-1.93% (7)
10	-9.59% (4)	-4.76% (5)	-3.51% (7)	0.58% (9)
15	-7.87% (6)	-0.81% (6)	-1.43% (7)	4.69% (8)
30	-8.02% (5)	-8.08% (5)	-4.43% (6)	-7.03% (6)

4.2.6. Enhanced Trials - Predicting Directional Strength

In this series of experiments, regression models were created to predict the Directional Strength of American Airlines stock. The configuration of the Proposed Method involved 4 communities, with each community comprising 100 agents. Three inputs were used to represent the Technical Indicators, while 5 outputs were utilized. To observe the method's behavior with a higher number of agents, a larger agent count was employed.

The experiments were conducted separately, with the first one running for 10 generations, the second for 100 generations, and the third for 1000 generations. The primary objective of these experiments was to demonstrate how an increase in the number of generations could lead to a reduction in prediction error. The objective was successfully achieved,

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Table 8
Microsoft Rates of Return

	TRAINING SET SIZE			
GEN	2	4	6	8
5	8.35% (10)	12.26% (11)	6.19% (9)	-1.52% (8)
10	11.82% (10)	12.26% (11)	1.48% (7)	6.97% (9)
15	8.35% (10)	12.26% (11)	11.74% (10)	6.97% (9)
30	8.35% (10)	12.26% (11)	11.74% (10)	2.42% (9)

Table 9
Average Rate of Return for Each of the 30 Stocks

Stock	Average Rate of Return	Stock	Average Rate of Return
AA	-0.917899%	AXP	-0.0986292%
BA	1.230233%	BAC	9.67739%
CAT	-0.0291421%	CSCO	1.223488%
CVX	-1.715249%	DD	1.653839%
DIS	-0.284699%	GE	-1.2747544%
HD	2.838411%	HPQ	-1.246380%
IBM	6.06691%	INTC	1.29393%
JNJ	-2.636135%	JPM	-0.816559%
KO	2.293977%	KRFT	-0.353017%
MCD	-2.848011%	MMM	-0.35037%
MRK	-6.39179%	MSFT	-4.97989%
PFE	0.934639%	PG	-2.91554%
T	0.0619734%	TRV	-1.40226%
UTX	1.27280%	VZ	-2.74370%
WMT	-2.33238%	XOM	1.87789%

as shown in Table 10. The corresponding curve-fitting charts for each experiment can be found in Figure 10 (10 generations), Figure 11 (100 generations), and Figure 12 (1000 generations).

4.3. Contribution of our work

our contributes to the existing field of knowledge and extends recent developments in the following ways

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Table 10
Directional Strength Prediction

Generations	abs(MSE + sum of forces)
10	4430.62
100	2573.73
1000	2310.74

- Holistic Modeling: Our method presents a holistic modeling framework considering financial market technical, fundamental, and emergent factors. This comprehensive approach builds on recent hybrid models, providing a more nuanced understanding of market behavior.
- Agent-Based Explanation: By employing agent-based modeling, our work extends the capability of financial forecasting models to simulate agent interactions and market dynamics. This advancement contributes to a deeper exploration of market forces and behaviors.
- Innovative Algorithm Fusion: Our method's fusion of fuzzy inference systems, genetic programming, and multi-agent systems is innovative and unique. This integration improves forecasting accuracy and extends our model's range of scenarios.
- Practical Applicability: Our work focuses on decision support system (DSS) development, enhancing its practical applicability. By recommending trading strategies, your method bridges the gap between theoretical models and actionable insights for traders.

In summary, our work builds upon recent trends in hybrid models, deep learning, ensemble methods, and explainability while introducing novel elements such as agent-based modeling and innovative algorithm fusion. These contributions advance the state of the art by providing a holistic, interpretable, and practically applicable approach to financial forecasting.

5. Conclusion and Future Work

In summary, our research unveils a robust and forward-looking decision support system with significant promise for shaping stock selection and redefining investment strategies. The unique amalgamation of a multi-agent paradigm, fuzzy inference systems, and genetic programming positions this method at the vanguard of data-driven financial analysis. The system showcases an impressive ability to forecast stock movements with heightened accuracy through the dynamic adaptation of parameters and the precision of regression tasks. Additionally, we intend to modify GP algorithm operators to improve the membership functions and dynamically adjust literal ranges.

While acknowledging the necessity for ongoing validation and refinement, we envisage our approach as a potent tool that empowers investors and financial analysts alike, delivering actionable insights amidst the complexity of modern financial markets. As our endeavors persist in refining and substantiating this approach, we anticipate its tangible impact on data-driven financial decision-making, ushering in a new era of informed choices and strategic advancements

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

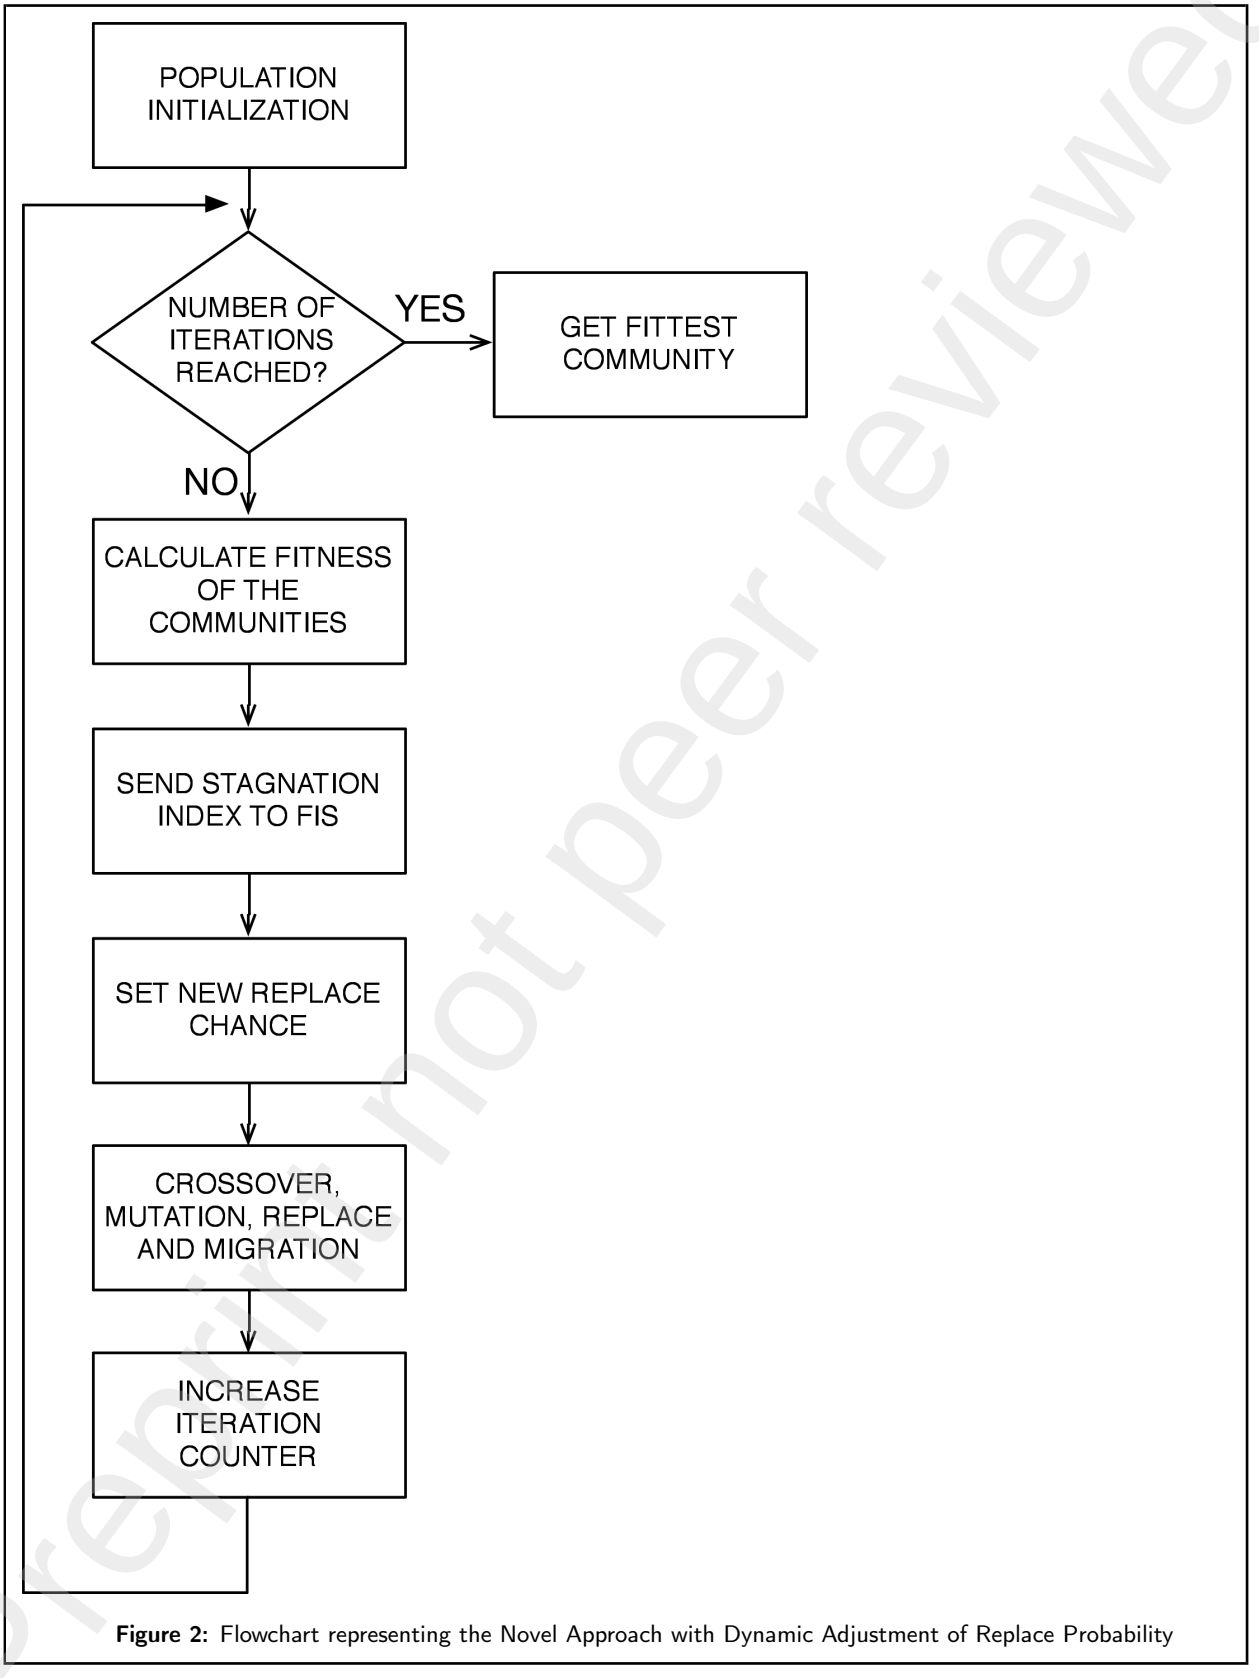

Figure 2: Flowchart representing the Novel Approach with Dynamic Adjustment of Replace Probability

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

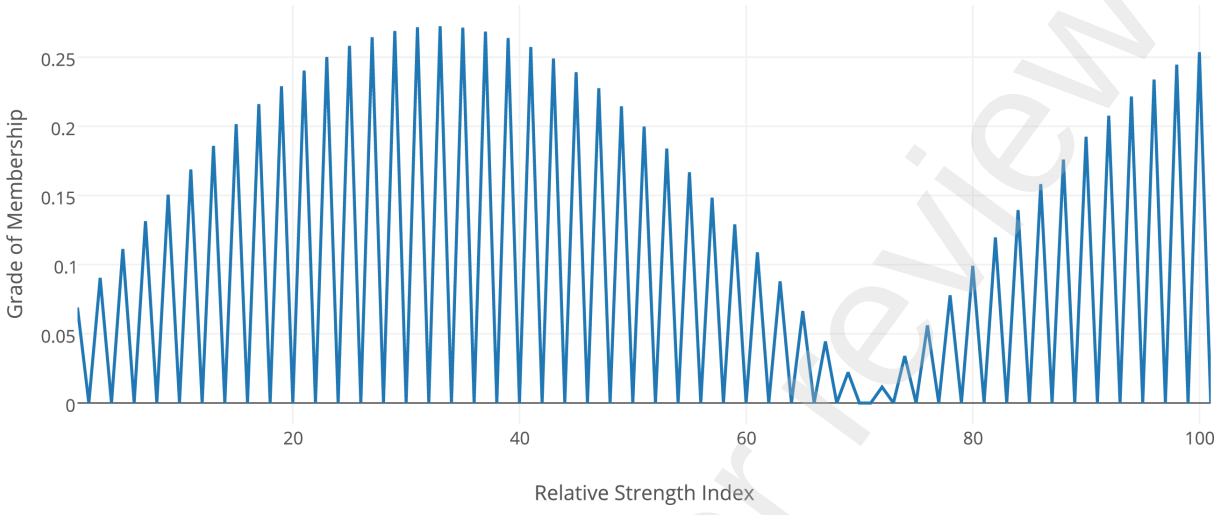

Figure 3: Visualization of Membership Function for the Input Variable

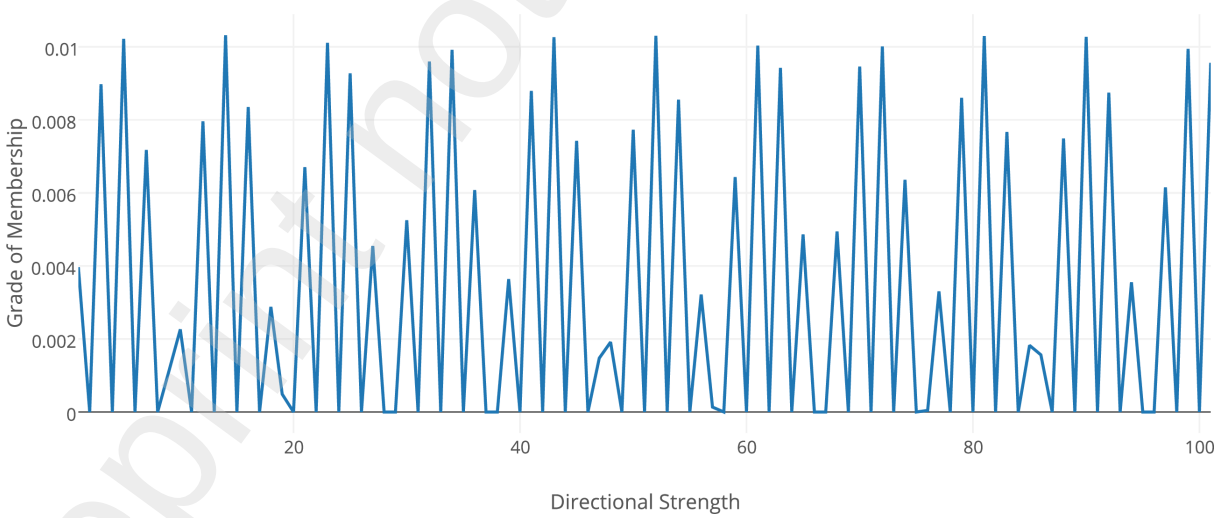

Figure 4: Visualization of Membership Function for the Output Variable

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

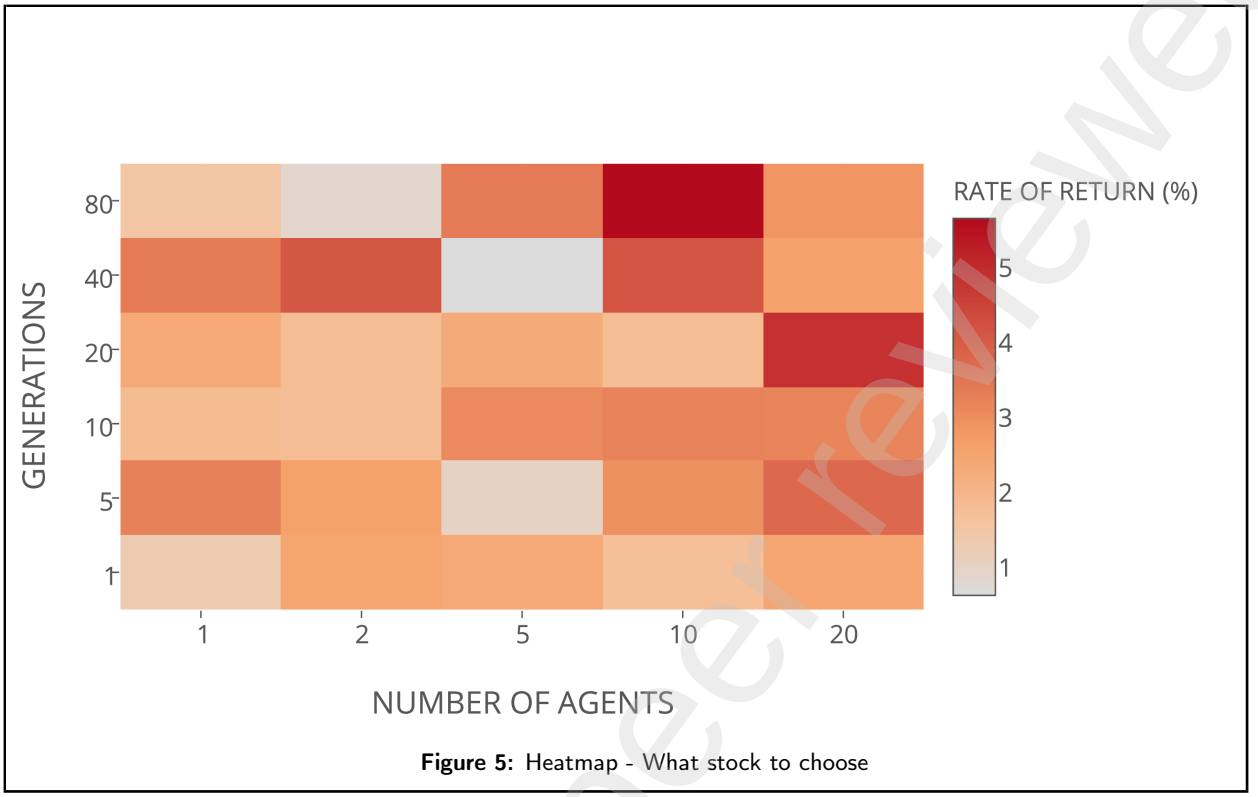

Figure 5: Heatmap - What stock to choose

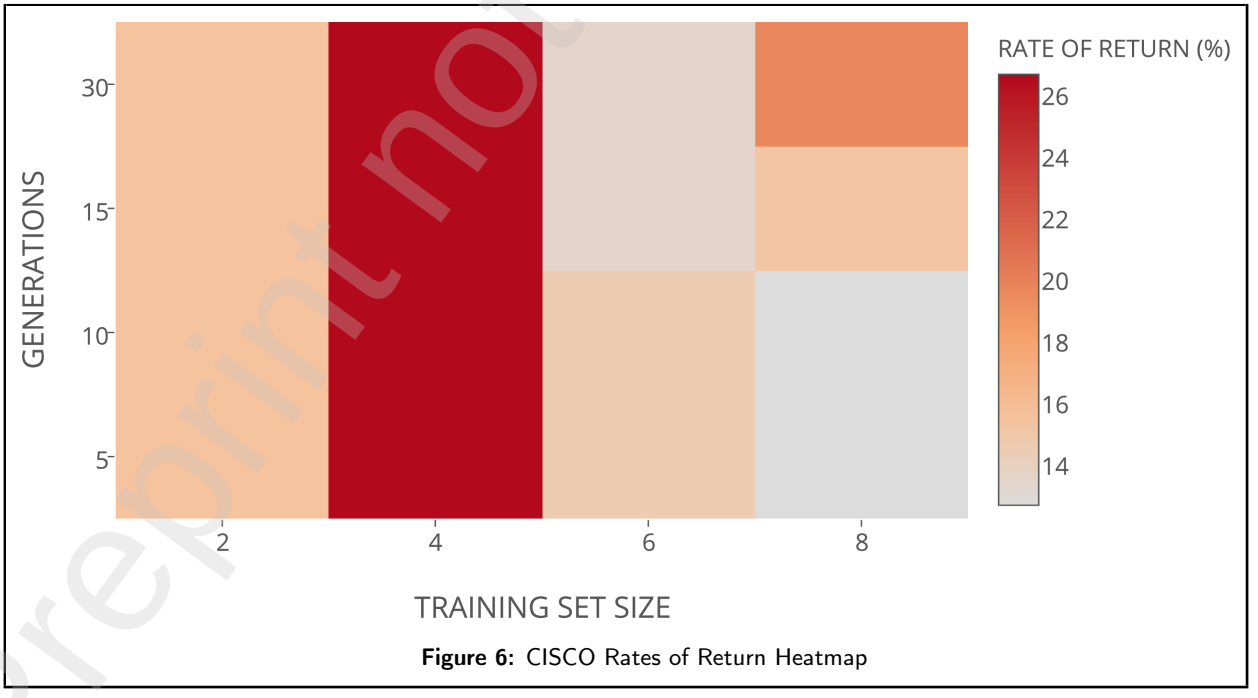

Figure 6: CISCO Rates of Return Heatmap

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

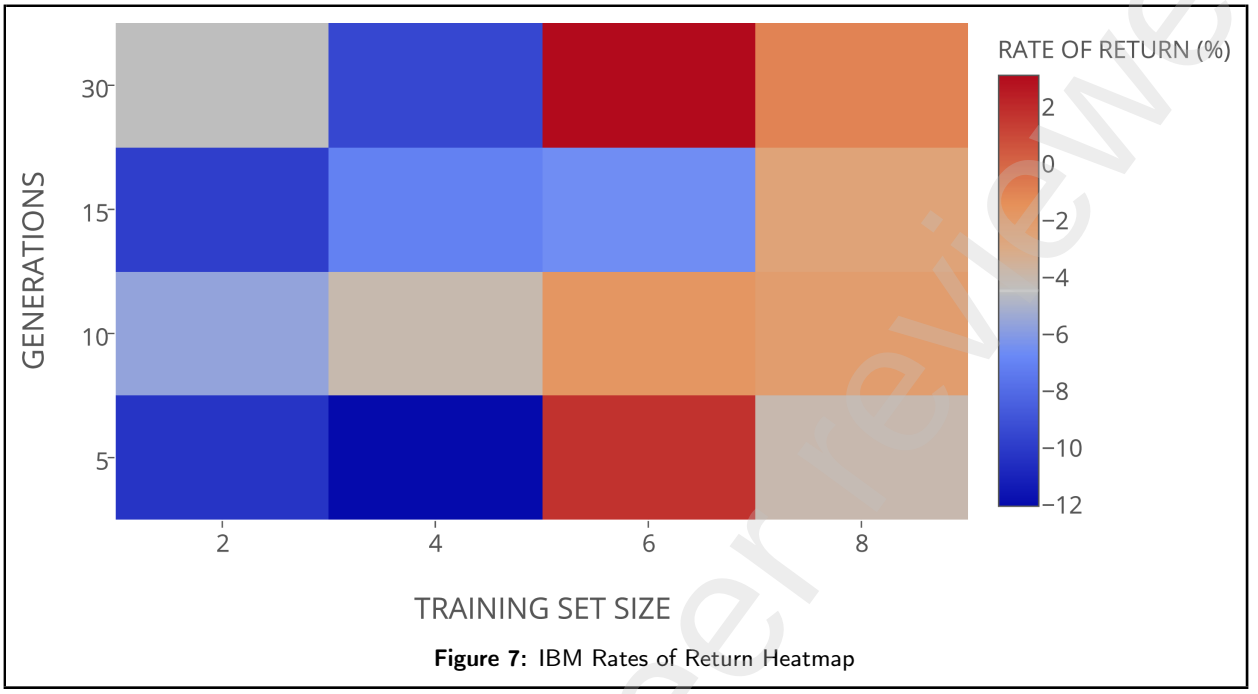

Figure 7: IBM Rates of Return Heatmap

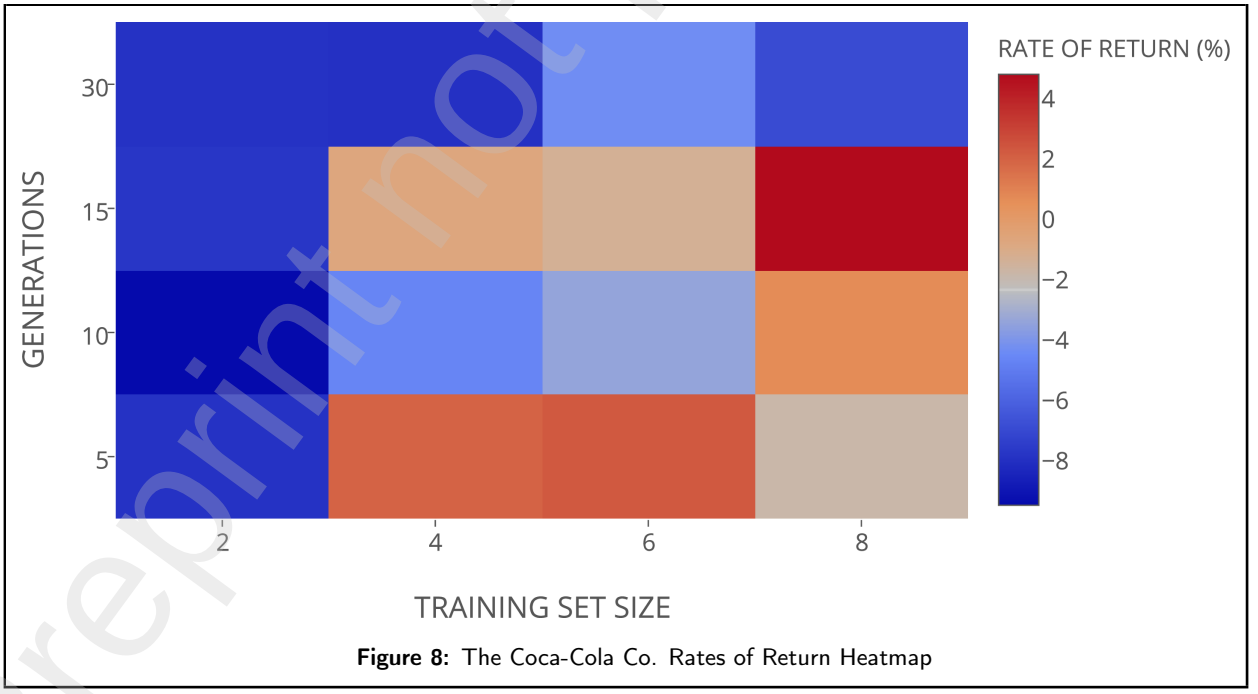

Figure 8: The Coca-Cola Co. Rates of Return Heatmap

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

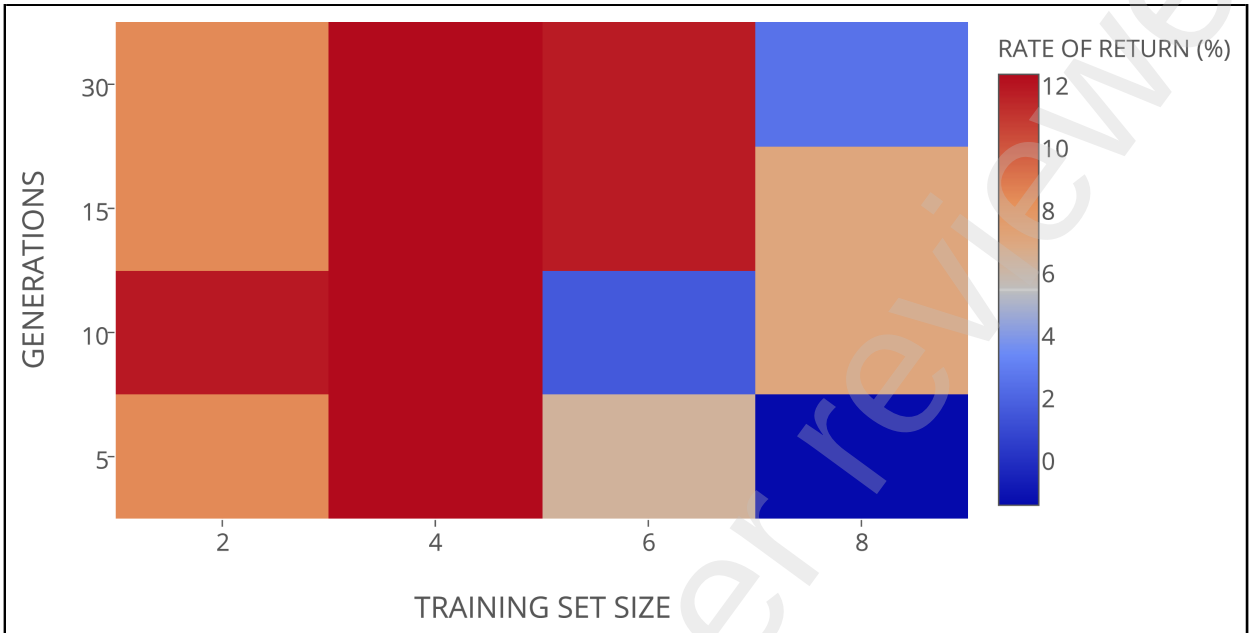

Figure 9: Microsoft Rates of Return Heatmap

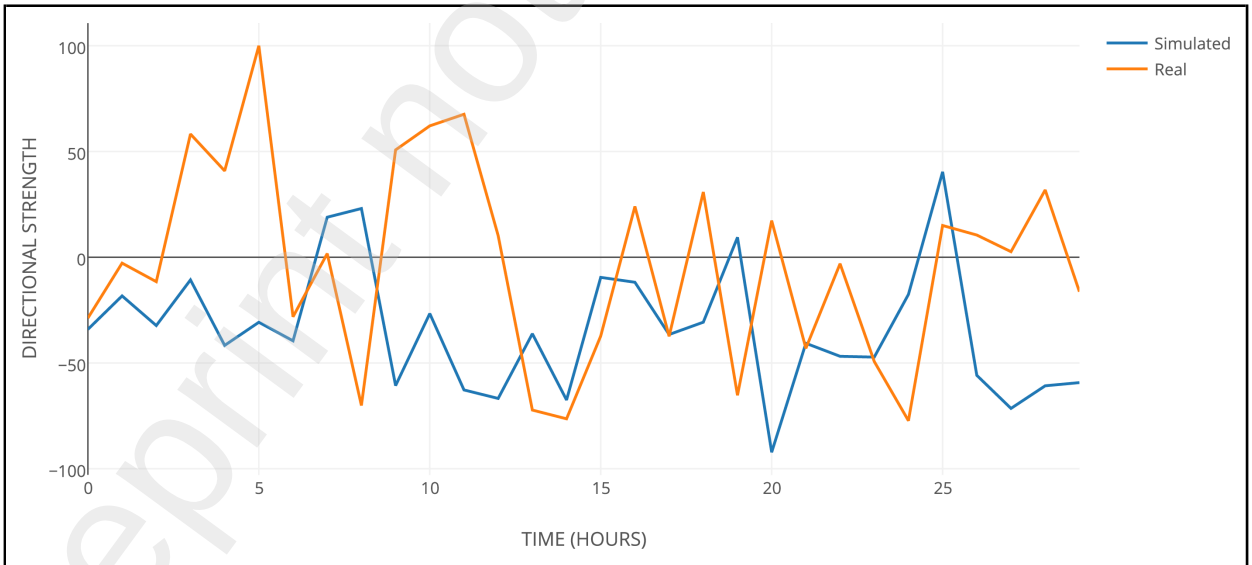

Figure 10: Directional Strength Simulation - 100 Agents, 10 Generations

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

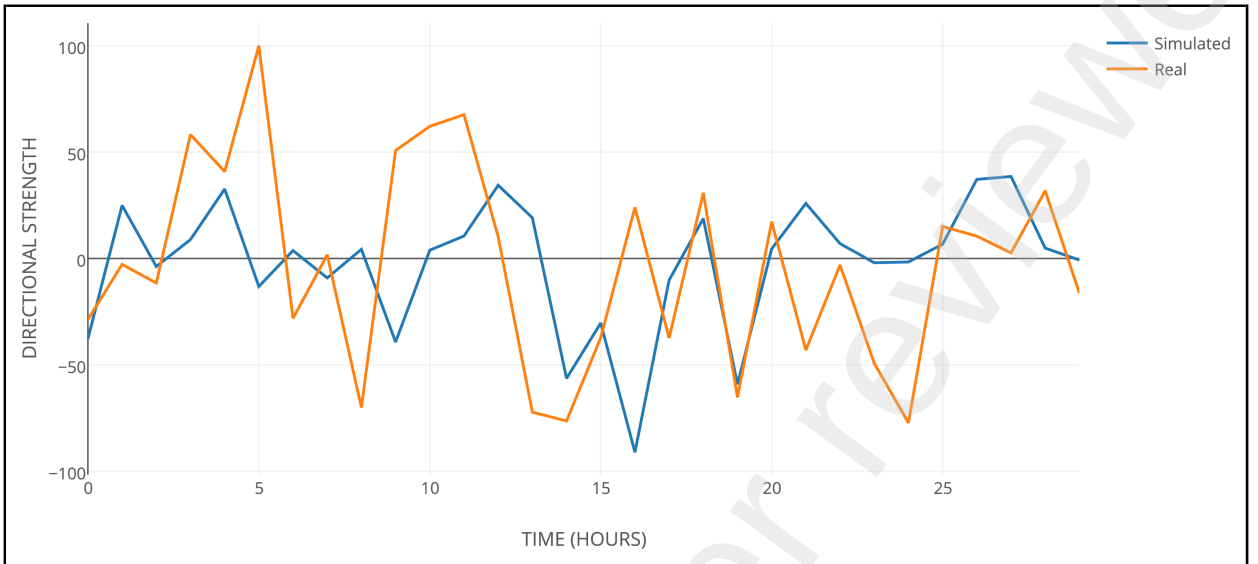

Figure 11: Directional Strength Simulation - 100 Agents, 100 Generations

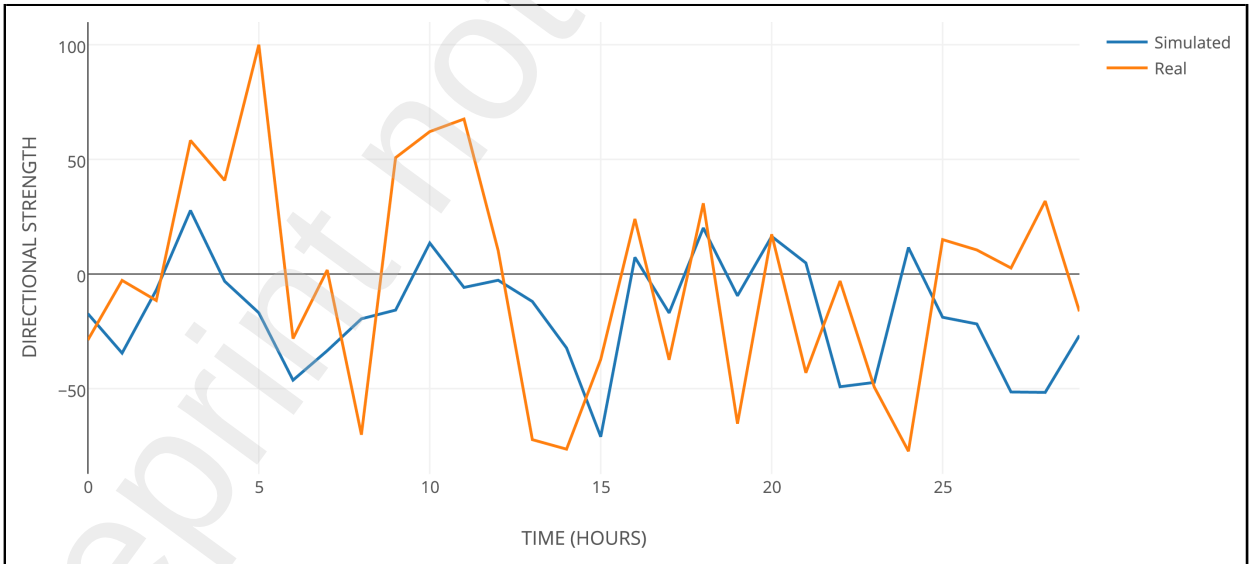

Figure 12: Directional Strength Simulation - 100 Agents, 1000 Generations

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

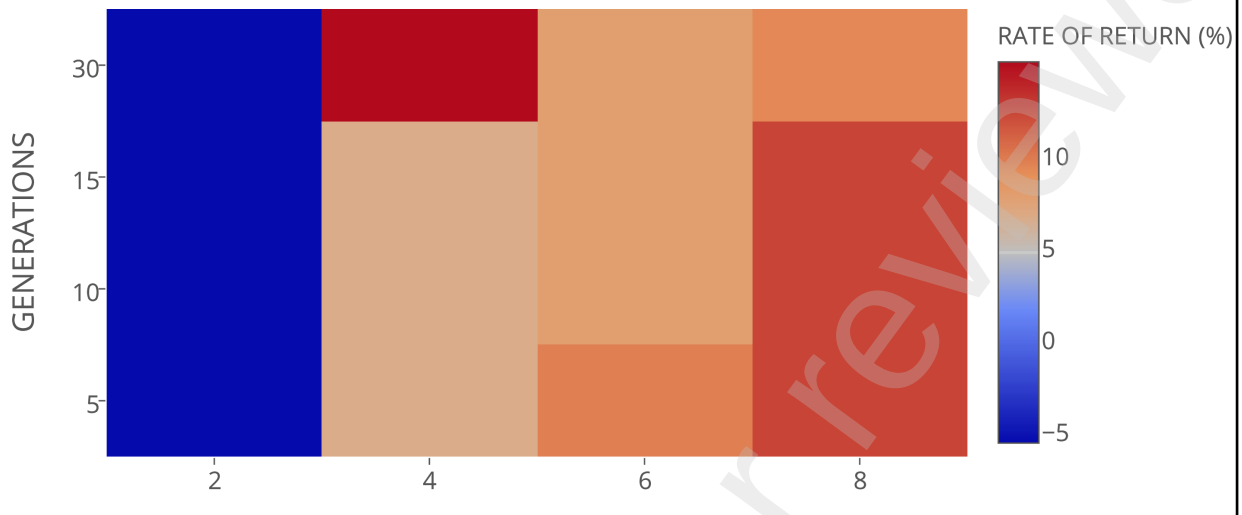

Figure 13: American Airlines Rates of Return Heatmap

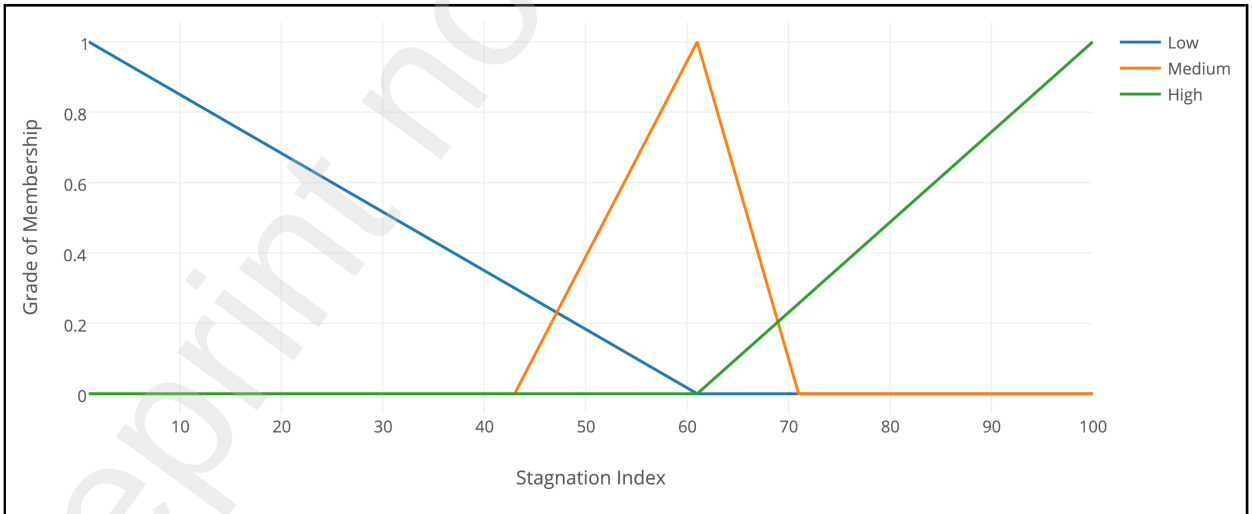

Figure 14: Input for the Dynamic Adaptation of the Replace Chance

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

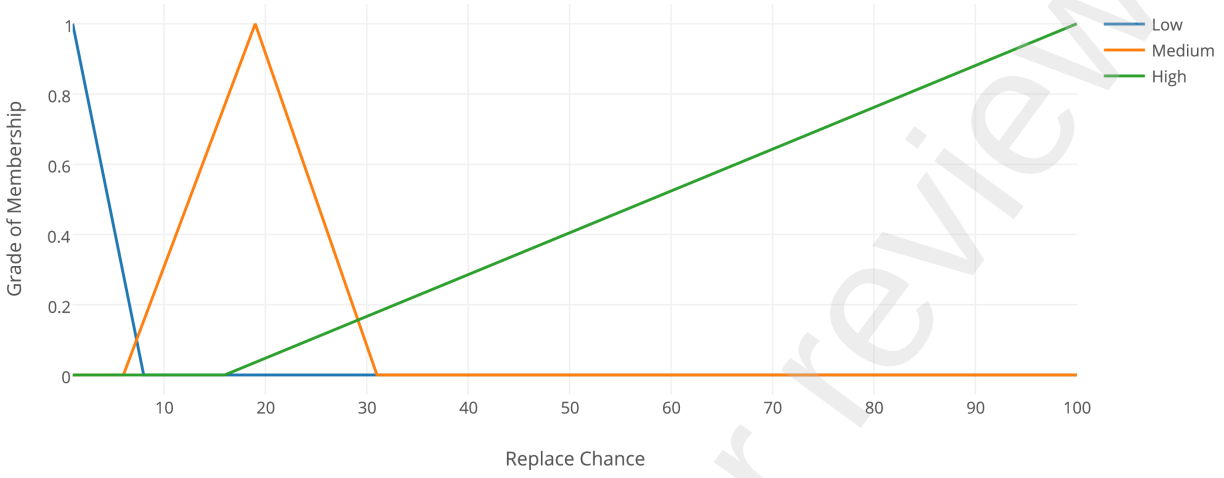

Figure 15: Output for the Dynamic Adaptation of the Replace Chance

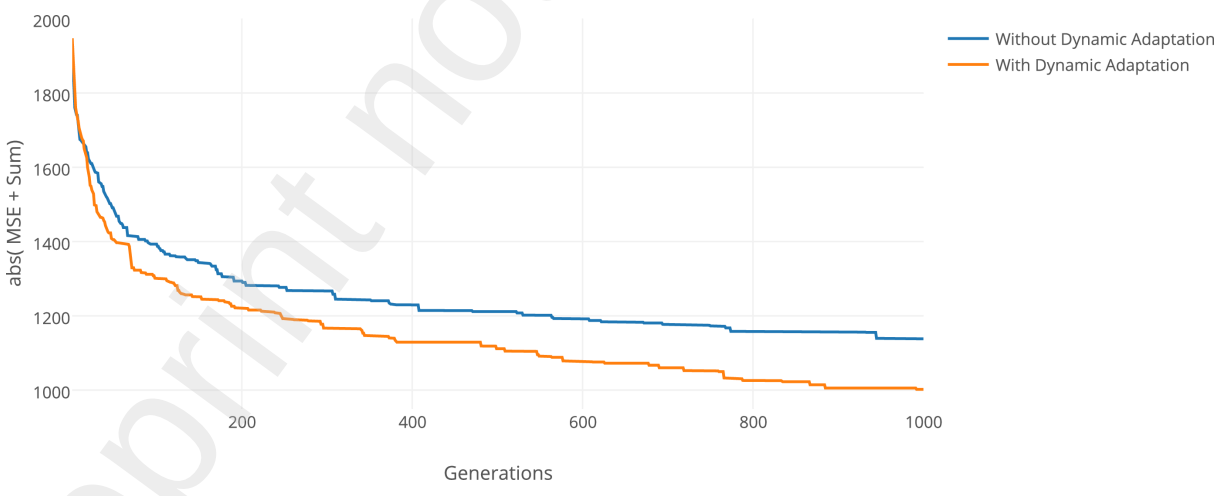

Figure 16: Dynamic Adaptation of Replace Chance vs Constant Replace Chance

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Two-sample t test with equal variances

Variable	Obs	Mean	Std. Err.	Std. Dev.	[99% Conf.	Interval]
with_fis	30	1142.388	20.95996	114.8024	1084.614	1200.162
withou~s	30	1244.283	22.00501	120.5264	1183.629	1304.937
combined	60	1193.335	16.46105	127.5067	1149.52	1237.151
diff		-101.8954	30.38981		-182.8322	-20.95866

diff = mean(with_fis) - mean(without_fis) t = -3.3529

Ho: diff = 0 degrees of freedom = 58

Ha: diff < 0 Ha: diff != 0 Ha: diff > 0

Pr(T < t) = 0.0007 Pr(|T| > |t|) = 0.0014 Pr(T > t) = 0.9993

Figure 17: Hypothesis Test for the Dynamic Adaptation of the Replace Chance

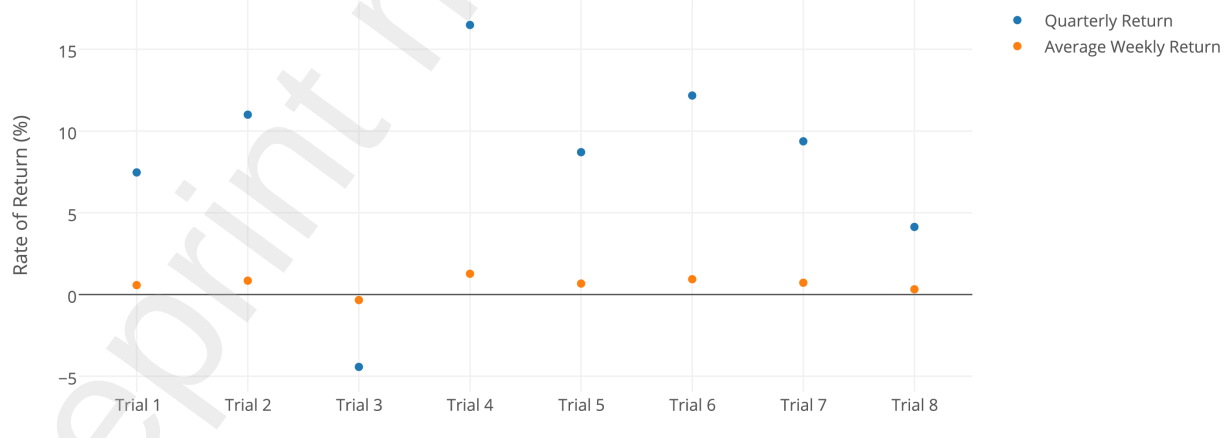

Figure 18: Rate of Return of the 8 Trials

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

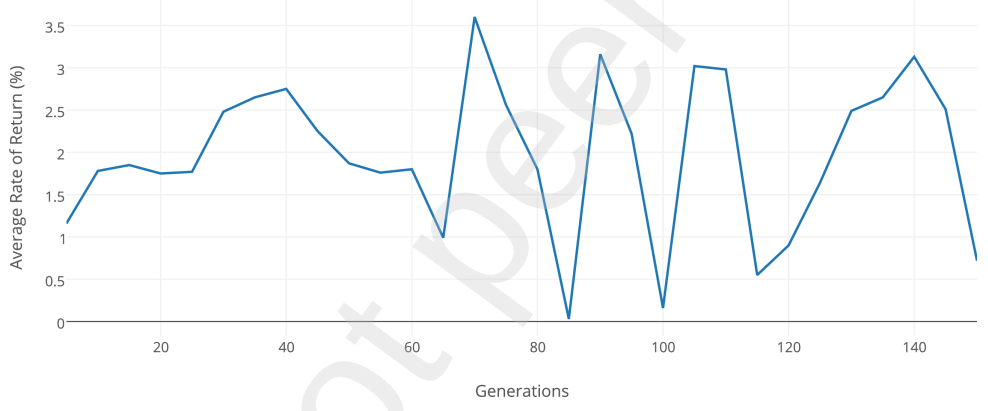

Figure 19: Experiments of the Proposed Method with a Constant Configuration and Increasing Generations

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

References

Alam, K.S., Lee, S.G., Egiazarian, K., 2018. Stock price prediction using lstm, rnn and cnn-sliding window model. Neurocomputing 321, 321–331.

Basak, D., Pal, S., Patranabis, D.C., 2007. Support vector regression. Neural Information Processing-Letters and Reviews 11, 203–224.

Bastian, A., 2000. Identifying fuzzy models utilizing genetic programming. Fuzzy Sets and Systems 113, 333–350. doi:10.1016/S0165-0114(98)00086-4.

Billard, L., Diday, E., 2002. Symbolic regression analysis, in: Classification, Clustering, and Data Analysis. Springer, pp. 281–288.

Boer-Sorban, K., 2008. Agent-Based Simulation Agent-Based Simulation of Financial Markets - A modular, Continous-Time Approach. Romania .

Bonabeau, E., 2002. Agent-based modeling: Methods and techniques for simulating human systems. Proceedings of the National Academy of Sciences 99, 7280–7287.

Brown, M.S., Pelosi, M.J., Dirska, H., 2013. Dynamic-radius species-conserving genetic algorithm for the financial forecasting of dow jones index stocks, in: Machine Learning and Data Mining in Pattern Recognition. Springer, pp. 27–41.

Burg, J.P., 1968. A new analysis technique for time series data. NATO advanced study institute on signal processing with emphasis on underwater acoustics 1.

Castillo, O., Melin, P., 2001. Simulation and forecasting complex economic time series using neural networks and fuzzy logic, in: Neural Networks, 2001. Proceedings. IJCNN'01. International Joint Conference on, IEEE. pp. 1805–1810.

Castillo, O., Neyoy, H., Soria, J., Melin, P., Valdez, F., 2015. A new approach for dynamic fuzzy logic parameter tuning in Ant Colony Optimization and its application in fuzzy control of a mobile robot. Applied Soft Computing 28, 150–159.

Chen, S.H., Yeh, C.H., 2001. Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market. Journal of Economic Dynamics and Control 25, 363–393. doi:10.1016/S0165-1889(00)00030-0.

Connor, N.O., Madden, M.G., 2005. A Neural Network Approach to Predicting Stock Exchange Movements using External Factors. Applications and Innovations in Intelligent Systems XIII. SGAI 2005. Springer, London .

Dash, G., Behera, H.S., 2020. Explainable artificial intelligence for financial time series analysis: A survey. Expert Systems with Applications 141, 112920.

Fang, L., Wang, Y., Guo, S., Zhang, L., Wei, W., 2021. Hybrid deep learning model for stock price prediction using numerical and textual information. Neurocomputing 434, 328–337. doi:10.1016/j.neucom.2021.02.125.

Ferreira, D.F., da Costa, R.S., Rodrigues, R.J., Papa, J.P., 2019. Financial time series forecasting using stacked ensemble learning: A comprehensive review. Information Fusion 52, 176–194.

Gamil, A.a., El-fouly, R.S., Darwish, N.M., 2007. Stock Technical Analysis using Multi Agent and Fuzzy Logic. Proceedings of the World Congress on Engineering I, 6.

Garcia-Almanza, A.L., Tsang, E.P.K., 2006. Forecasting stock prices using Genetic Programming and Chance Discovery. 12th International Conference On Computing In Economics And Finance , number 489URL: http://ideas.repec.org/p/sce/scecfa/489.html.

Gilbert, G.N., 2008. Agent-based models. 153, Sage.

Grothmann, R., 2002. Multi-agent market modeling based on neural networks. Faculty of Economics, University of Bremen URL: http://lsc.fie.umich.mx/~juan/Materias/Cursos/ANN/Papers/multi-agent-market-modeling.pdf.

Homaifar, a., McCormick, E., 1995. Simultaneous design of membership functions and rule sets for fuzzy controllers using genetic algorithms. Fuzzy Systems, IEEE Transactions ... 3, 129–139. URL: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=388168, doi:10.1109/91.388168.

Huang, H., Pasquier, M., Quek, C., 2009. Financial market trading system with a hierarchical coevolutionary fuzzy predictive model. IEEE Transactions on Evolutionary Computation 13, 56–70. doi:10.1109/TEVC.2008.911682.

Hussain, M., Muhammad, T., Khan, I., 2019. Hybrid approaches for stock price forecasting: A survey. Expert Systems with Applications 129, 208–222.

Ijegwa, A.D., Rebecca, V.O., Olusegun, F., Isaac, O.O., 2014. A Predictive Stock Market Technical Analysis Using Fuzzy Logic. Computer and Information Science 7, 1–17. doi:10.5539/cis.v7n3p1.

Jennings, N.R., 2001. An agent-based approach for building complex software systems. Communications of the ACM 44, 35–41.

Keen, P.G., 1980. Decision support systems: a research perspective. Decision Support Systems: Issues and Challenges (New York: Pergamon Press, 1980) , 23–44.

Kendall, G., Su, Y., 2003. The co-evolution of trading strategies in a multi-agent based simulated stock market through the integration of individual learning and social learning. Proceedings of the 2003 Congress on Evolutionary Computation , 200–206URL: http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.62.9288&rep=rep1&type=pdf.

Klingert, F.M.A., Meyer, M., 2012. Comparing Prediction Market Mechanisms Using An Experiment-Based Multi-Agent Simulation, in: ECMS 2012 Proceedings edited by: K. G. Troitzsch M. Moehring, U. Lotzmann, ECMS. URL: http://dx.doi.org/10.7148/2012-0654-0661, doi:10.7148/2012-0654-0661.

Koza, J.R., 1992. The genetic programming paradigm: Genetically breeding populations of computer programs to solve problems. Dynamic, Genetic, and Chaotic Programming , 203–321URL: http://www.genetic-programming.com/jkpdf/soucek1992.pdf.

Kutner, M.H., Nachtsheim, C., Neter, J., 2004. Applied linear regression models. McGraw-Hill/Irwin.

Li, J., Tsang, E., 1999. Improving technical analysis predictions: an application of genetic programming. Proceedings of The 12th International Florida AI Research Society Conference, Orlando, Florida , 108–112URL: http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Improving+Technical+Analysis+Predictions+:+An+Application+of+Genetic+Programming#0.

Little, J.D., 1979. Decision support systems for marketing managers. The Journal of Marketing , 9–26.

Lu, C.J., Lee, T.S., Chiu, C.C., 2009. Financial time series forecasting using independent component analysis and support vector regression. Decision Support Systems 47, 115–125. URL: http://dx.doi.org/10.1016/j.dss.2009.02.001, doi:10.1016/j.dss.2009.02.001.

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

Martinez-Jaramillo, S., Tsang, E.P.K., 2009. An heterogeneous, endogenous and coevolutionary GP-based financial market. IEEE Transactions on Evolutionary Computation 13, 33–55. doi:10.1109/TEVC.2008.2011401.

Melin, P., Mancilla, A., Lopez, M., Mendoza, O., 2007. A hybrid modular neural network architecture with fuzzy Sugeno integration for time series forecasting. Applied Soft Computing 7, 1217–1226.

Nguyen, H.D., Nguyen, V.C., Nguyen, Q.V., 2019. Combining neural networks and support vector regression for stock price direction prediction. Expert Systems with Applications 135, 15–23.

Niazi, M., Hussain, A., 2011. Agent-based computing from multi-agent systems to agent-based models: A visual survey. Scientometrics 89, 479–499. doi:10.1007/s11192-011-0468-9.

Pellizzari, P., Forno, A.D., 2007. A comparison of different trading protocols in an agent-based market. Journal of Economic Interaction and Coordination 2, 27–43. doi:10.1007/s11403-006-0016-5.

Poli, R., Langdon, W.B., McPhee, N.F., Koza, J.R., 2008. A field guide to genetic programming. Lulu. com.

Russell, S.J., Norvig, P., Canny, J.F., Malik, J.M., Edwards, D.D., 2003. Artificial intelligence: a modern approach. volume 2. Prentice hall Upper Saddle River.

Samanidou, E., Zschischang, E., Stauffer, D., Lux, T., 2007. Agent-based models of financial markets. Rep. Prog. Phys. 70, 409–450. URL: http://dx.doi.org/10.1088/0034-4885/70/3/r03, doi:10.1088/0034-4885/70/3/r03.

Schirding, H., 1984. Stochastic oscillator. Technical Analysis of Stocks and Commodities 3, 97.

Sherstov, a., Stone, P., 2005. Three automated stock-trading agents: A comparative study. Agent-Mediated Electronic Commerce VI. Theories ... 18, 433–41. URL: http://www.springerlink.com/index/1n177848nj158772.pdf$\backslash$nhttp://www.ncbi.nlm.nih.gov/pubmed/8364599$\backslash$nhttp://ukpmc.ac.uk/abstract/MED/8364599, doi:10.1007/11575726_13.

Shoham, Y., Leyton-brown, K., 2009. Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. ReVision doi:10.1073/pnas.97.18.9840.

Sprague, R.H., 1980. A Framework for the Development of Decision Support Systems. MIS Quarterly 4, 1–26. URL: http://www.jstor.org/stable/248957, doi:10.1109/ICMIT.2008.4654367.

Thrift, P., 1991. Fuzzy Logic Synthesis with Genetic Algorithms.

Tsang, E., 2008. Forecasting – where computational intelligence meets the stock market The bad news : the efficient market hypothesis. https://doi.org/10.1007/s40745-022-00432-6 .

Tsang, E., Yung, P., Li, J., 2004. EDDIE-automation, a decision support tool for financial forecasting. Decision Support Systems 37, 559–565. doi:10.1016/S0167-9236(03)00087-3.

Tsang, E.P., Li, J., Butler, J.M., 1998. EDDIE beats the bookies. Softw., Pract. Exper. 28, 1033–1043.

Tsang, E.P., Li, J., Markose, S., Er, H., Salhi, A., Iori, G., 2000. EDDIE in financial decision making. Journal of Management and Economics 4.

Valdez, F., Melin, P., Castillo, O., 2014. A survey on nature-inspired optimization algorithms with fuzzy logic for dynamic parameter adaptation. Expert Systems with Applications 41, 6459–6466.

Wilder, J.W., 1978. New concepts in technical trading systems. Trend Research.

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132

CogniTrade: A Multi-Agent Framework for Data-Driven Financial Decision Making Using Fuzzy Inference and Genetic Programming

Srikari Rallabandi[1] and Manleen Kaur[2] Pranav Misra

[1] Vidya Jyothi Institute of Technology, Hyderabad India
sreeku.ralla@gmail.com
[2] Kalinga University, Raipur, India
manlinsarna74@gmail.com
[3] Dr. A.P.J Abdul Kalam Technical University, India
pranavmisra7@gmail.com

Abstract. This paper introduces a pioneering decision support system that revolutionizes stock selection through an innovative blend of a multi-agent paradigm, fuzzy inference systems, and genetic programming. In this framework, a collaborative network of agents harnesses fuzzy logic systems to decipher optimal trading directions and strengths, shedding light on the intricate nuances of financial market behavior. To arrive at a decision, an architecture of a zero-sum multi-agent system based on fuzzy inference systems and genetic programming was designed. The system adeptly predicts directional movements and magnitudes by conducting regression tasks on stock time series data. The multi-agent system performs a regression task on the time series data of a stock to determine the direction and strength of the next movement in the time series; By adopting this framework, individuals can examine its affiliation functions to comprehend financial market behavior better. Rigorous testing against a Dow Jones Index benchmark underscores the method's exceptional profitability, outperforming established techniques. With its potential to amplify investment decisions through insightful data analysis, this system emerges as a transformative force in financial decision-making.

Keywords: Artificial intelligence, Financial forecasting, Fuzzy Systems, Genetic Programming, Multi-Agent Systems

Preprint not peer reviewed

This preprint research paper has not been peer reviewed. Electronic copy available at: https://ssrn.com/abstract=4562132